IAF-99-B.6.01
SAFS: THE STANDARD AUTONOMOUS FILE SERVER

Susan K. Semancik
NASA Goddard Space Flight Center's Wallops Flight Facility

ABSTRACT

The purpose of the Standard Autonomous File Server (SAFS) is to provide automated management of large data files without interfering with the assets involved in the acquisition of the data. It operates as a stand-alone solution, monitoring itself, and providing an automated level of fail-over processing to enhance reliability. By using an improved automated file transfer process, the SAFS system provides a quicker, more reliable file distribution for customers of near real-time data than has been realized by previous methods. In addition, when overlapping supports occur in multiple projects, provision is made for file transfer prioritization by bandwidth sharing or transfer interruption methods. Web reporting provides current status of system availability, file latency, and customer file distribution

Initially, the SAFS system has been installed at NASA Ground Network sites for distributed acquisition of satellite data in support of QuikSCAT and ADEOS II missions. It has also been installed at NASA Goddard Space Flight Center to provide for centralized customer data distribution. Since the Summer of 1998, these SAFS systems have been successfully involved in the integration, testing, and launch support events with the QuikSCAT Mission Operations Center at the University of Colorado, NASA ground stations, and QuikSCAT customers. SAFS is proving to be both timely, and dependable, and should add significant value to the science data community. The purpose of this document is to describe the mission requirements, the design and functional specifications, and operations concept underlying this initial application.

IAF-99-B.6.01

SAFS: THE STANDARD AUTONOMOUS FILE SERVER

Susan K. Semancik
NASA Goddard Space Flight Center's Wallops Flight Facility
Wallops Island, Virginia USA

50th International Astronautical Congress
4-8 Oct 1999/Amsterdam, The Netherlands

SAFS: THE STANDARD AUTONOMOUS FILE SERVER

Susan K. Semancik
NASA Goddard Space Flight Center's Wallops Flight Facility

1.0 Introduction

1.1Purpose

The purpose of the Standard Autonomous File Server (SAFS) is to provide automated management of large data files while providing timely access to these files without interfering with the assets involved in their acquisition. The broad operational concept of the SAFS system allows for its placement in diverse environments meeting a variety of needs. For example, the SAFS system could have merit in an operational launch environment to give experimenters, operations personnel, educators, and others access to raw vehicle data without interfering with the operation of the range. The SAFS system could be a standard resource in any large transaction based system in which the data source is real time and time constraints and/or security dictate non-interference by external access. The purpose of this document is to describe the mission requirements, the design and functional specifications, and operations concept underlying the SAFS effort in its first application.

1.2 Initial Application

The initial application of SAFS has been as an autonomous intermediary between the NASA Ground Networks (NGN) and their satellite data customers whose latency requirements cannot be met by media distribution. For SAFS purposes, file latency is defined to be the time from the start of satellite data downlink to the availability of the file to the customer. Telemetry processors at the ground stations acquire raw data from these downlinks and provide data files to the SAFS system for later customer consumption.

The first NGN projects SAFS was designed to support are QuikSCAT, launched June 19, 1999, and ADEOS-II, expected to be launched in late 2000. These projects have multiple file types with associated metadata files, acquired in up to 14 passes a day distributed among ground stations at Svalbard, Norway (SGS); Poker Flat, Alaska (AGS); Wallops Orbital Tracking Station, Virginia (WGS); Alaska SAR Facility, Alaska (ASF); and McMurdo, Antarctica (MGS). The file sizes range from 3 MB to 370 MB, and have multiple customers with file latencies from 1 to 11 hours. Interest has been strong from other projects desiring SAFS support, and it is likely that two to five additional projects will be supported before the end of 1999.

2.0 Design Strategy

2.1 Network Environment

The network environment in which SAFS is placed is crucial to meeting customer latency requirements. While SAFS can not control the external network components, it does use, wherever possible, an improved file transfer protocol to guarantee file delivery with self-correcting transmissions, recovery from point of failure, and options for stop/resume transmission control, on-the-fly compression, and programmable network bandwidth control. This is accomplished through the use of a COTS product, FASTCopy (www.softlinkusa.com), which also provides an application programming interface for automated tasks, pre-transfer processing, post-transfer processing, and scheduled transfers.

2.2 Operational Concept

The operational design for NGN support incorporates distributed SAFS systems at ground stations on closed networks for file acquisition from telemetry processors (TMP), and a centralized SAFS at NASA Goddard Space Flight Center (GSFC) on open networks for file distribution to project customers. The Central SAFS provides a single point of contact for customers and isolates the ground stations from customer interactions. At each ground station, multiple TMP's supporting multiple antennas and/or multiple projects, acquire downlinked satellite data that is processed into files and sent to the ground station SAFS as shown in Figure 1:

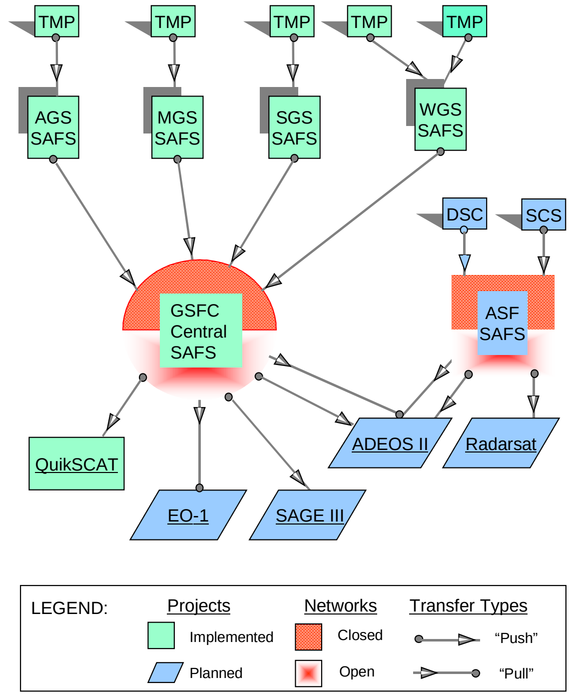

Figure 1: SAFS configuration as of mid 1999

Each of these SAFS systems uses FASTCopy to automatically push these files to the central SAFS, where they are made available to each project's customers.

Customers can "pull" a file from the SAFS system once they receive a data ready notification (DRN) of its availability. Or if they are also using the same COTS product, FASTCopy, the SAFS system can automatically "push" their files to them, thereby eliminating the delay inherent in the notification and reaction processes required for "pull" customers. "Push" customers are emailed a file delivery notice (FDN) with transfer status upon completion of the file transfer to them. Both FDN and DRN notices contain the file's name, size, and its SAFS location.

Each SAFS is a truly autonomous system and can be used independently of the Central SAFS configuration, as can be seen in Figure 1, by its deployment at the University of Alaska's SAR facility. Customers acquire files directly from the ASF SAFS, which is connected to multiple telemetry processors.

2.3 System Configuration

The components of the SAFS system hardware can be configured in a standard 19" rack, and include the following:

- RAID storage system
- RAID monitoring/configuration system
- server
- rack mounted keyboard/monitor/touch pad.

In addition, the SAFS system contains the following software components:

- COTS file transmission software
- Custom scripts for job control and monitoring
- Web reporting software

The hardware components were chosen for speed, reliability, expandability, and with cost in mind. The servers are SGI Origin 200 and 2000 systems, for the ground station and central SAFS, respectively. The RAID system is the DataDirect Networks EV1000 RAID system, which provides automatic failover for drives and controllers without the loss of data, and hot-swappable fans and power supplies which are able to be replaced without taking the system off-line. The RAID system also has an independent monitoring and control PC system, which frees the server from performing RAID configuration, diagnostics and status reporting. An integrated rack-mounted keyboard/monitor/touch pad mounts in only 3" of rack space, and is used to give access to both the PC monitoring system and the server. Different sized RAID drive configurations are based on each site's projected worst-case daily volume and project customer retention times. The FASTCopy COTS file transmission software was chosen for its guaranteed delivery, ability to do automatic pre- and post- transfer operations, bandwidth and priority control, and multi-platform support.

Failover contingency plans include redundant components and automatic monitoring with notifications. A heartbeat is sent to automated tracking station master systems to indicate operational status of the ground station SAFS.

3.0 Implementation

3.1 Schedule

The high level requirements for this first application include automation of file acquisitions, file transmissions, email notifications, disk management, and performance reporting. The SAFS system design incorporates three builds to meet its project requirements:

- Single project support with NGN ground station and central SAFS installations,
- Multiple project support with non-NGN ground station installation, prioritizing simultaneous file transfers, and expanded status reporting,
- Automation of project additions, changes, and deletions to both systems and reports, streaming of web data, and finalization of manuals

for system operation and management.

Each build includes the following phases: requirements analysis, design, procurement, development, system installation, and project integration and testing

The integration of each build after the first requires careful management so as not to disrupt the operations of the previous build. The first build was started in June 1997, and since June 1998, the SAFS systems have successfully supported integration, testing, and launch events with the QuikSCAT Mission Operations Center at the University of Colorado, NASA ground stations, and QuikSCAT customers. By the Fall of 1998, build 2 was started, and will be in testing by the fourth quarter of 1999, with build 3 scheduled for completion in 2000.

3.2 File Acquisition

Automation of file acquisitions is accomplished through the use of "hot" directories to detect when files arrive on the SAFS system. These directories have a monitoring process with a time-selectable interval for checking file arrivals. It is possible that a file transferred from a TMP might take longer than the checking interval for the "hot" directory. When this happens, only a partial file would be detected and processed. To avoid this problem, the "hot" directory can be set to react to the arrival of a flag file, which contains only the names of the files to process. The TMP sends the flag file last, after sending the data files to the "hot" directory. When the monitoring process detects the flag file, it triggers the moving of the named data files to a temporary directory for pre-processing by a job file and the deletion of the flag file. The flag file is not needed if FASTCopy is being used for file transfer, as it can be set to use a temporary directory to completely capture the file before automatically moving it to the "hot" directory for processing, or it can automatically start the processing in a non-hot directory. The flag file can also be used to process files in groups, such as a data file with an accompanying metadata file.

Each file type for a project has a job file which processes the file by customer lists, so the same file can trigger a DRN being sent to some customers, and a file transfer with an FDN to other customers. A post-process controls archiving these files for a project-specific time, usually 96 hours, and updates the system's event log accordingly.

3.3 File Transmission

After a customer either "pulls" or receives a "pushed" file from the SAFS system, they send back a receipt confirmation notification, RCN, with a success or error status. This information is added to the SAFS event log for the file, and is especially useful in the SAFS performance monitoring and latency reports for "pull" customers.

3.3.1 Failover Options: Upon failure to deliver a file to a "push" customer's system, an FDN email is sent to the customer with a failure status so that the customer can pull the file when ready during the project-specified retention period. Failover options, as shown in Figure 2, are available to

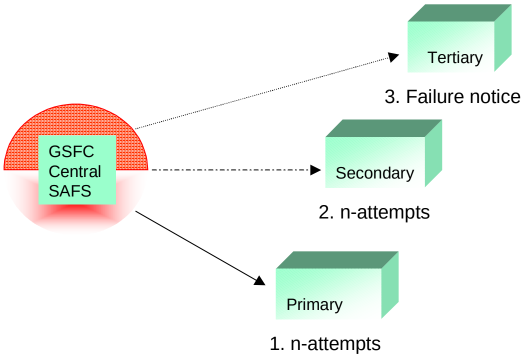

Figure 2: Failover "push" options to customer's systems.

"push" customers to circumvent delivery problems due to the unavailability of their system or network problems.

If a file cannot be delivered after three attempts, and the customer has indicated the availability of backup systems, then the file delivery is automatically attempted to their secondary system at least three times. Failure of the file transfer to the secondary system results in a failure FDN being sent to the customer's tertiary system, and the customer is expected to "pull" the file when ready.

3.3.2 Prioritization: The normal FIFO, First In First Out, scheme works fine for file transfers that don't overlap. But multi-project support increases the likelihood of simultaneous file arrivals and contention for transfer bandwidth. Even within a single project, contentions might occur if a large file is still being transferred when a higher priority file arrives. The network could be upgraded for additional bandwidth to account for the worst case for simultaneous transfers based on the total number of files and their sizes. But this potentially expensive solution should be considered only if satellite pass scheduling and file prioritization can not manage the distribution of the load to meet file latencies.

To handle simultaneous file arrivals and overlapping file transfers, the SAFS system combines bandwidth control with file prioritization. Four general prioritization levels have been defined to date:

- Special: for spacecraft or weather emergencies, or launch and early orbit activities
- High: for files needing speediest delivery
- Medium: for normal delivery
- Low: for least urgent delivery

There are categories within each priority level to allow for differences due to projects, file types, customer distributions, and latency requirements.

Before a file is transferred from a SAFS system, its priority is compared against that of any file already being transferred from the system. If the new file's priority is higher, then the transfer of the lower priority file is stopped, and the higher priority file is transferred. When it completes, the previously stopped transfer is resumed from the point where it had been stopped – NOT from the beginning of the file. If the new file's priority is equal to that of a file being transferred, then the new file can either be queued until the previous file's transmission has completed, or the files can be transferred together by sharing the bandwidth equally.

3.3.3 Bandwidth Control: If multiple files of various sizes arrive simultaneously and are queued for serial transmission, then the time for the smallest file to complete transfer to its destination from its arrival time would depend on its position in the queue, as can be seen in Figure 3.

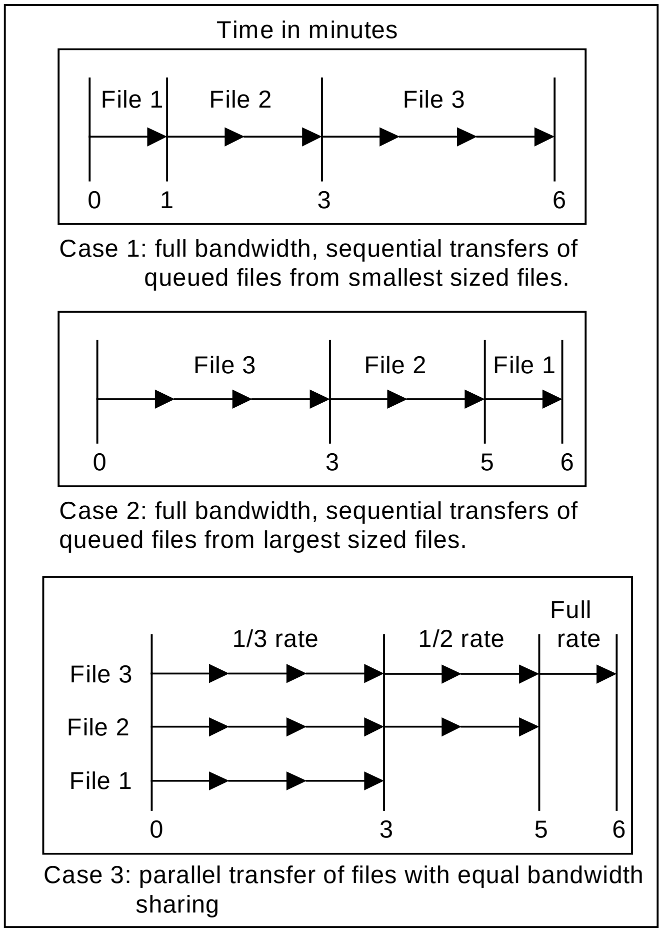

Case 1: full bandwidth, sequential transfers of queued files from smallest sized files.

Case 2: full bandwidth, sequential transfers of queued files from largest sized files.

Case 3: parallel transfer of files with equal bandwidth sharing

Figure 3: Sequential versus parallel transfers

If the files are transferred in parallel with equal bandwidth sharing, each file's transfer rate is slowed by increased network traffic and speeded up by a decrease in traffic. If n=number of simultaneous transfers, then 1/n = bandwidth available to each file, which results in a transfer rate of 1/n.

As can be seen in Table 1, the time for the smallest file to be serially transferred can range from fastest to slowest. Note that under the same network conditions, the time to transfer the files in parallel is the same as the total time to transfer them serially. And in parallel transfers, as shown in Table 1, the smallest file will arrive at its destination first. This is important in those projects that have a metadata file that needs to be examined first in order to decide on the use of its accompanying data file.

Files by size	Transfer time at full bandwidth (minutes)	Transfer time from arrival time (mins.)		
		case 1	case 2	case 3
File 1	I	1	6	3
File 2	2	3	5	5
File 3	3	6	3	6

Table 1: Summary of case transfer times for three simultaneously arriving files.

FASTCopy uses bandwidth sharing for files arriving simultaneously or which arrive while the SAFS is transferring other files of equal priority. The bandwidth can be set for a file transfer either absolutely or relatively, as a percentage of available bandwidth. Additionally, FASTCopy allows setting bandwidth with respect to time to allow scheduling for large file transfers.

These bandwidth sharing techniques are being utilized in the SAFS system, with job monitoring to avoid the following potential problems:

- A file never finishes transmitting if it keeps getting stopped by the arrival of higher priority files,
- The file transfers from one satellite pass are slowed so much by network traffic that they can not complete before the next pass of the same satellite.

3.4 Performance Reporting:

SAFS uses its own file naming convention internally, which contains timing and project information used in status reports derived from events associated with each file while on a SAFS system. These status reports are available on each SAFS system and are sent to a web server for presentation in one of two formats:

- An immediate "quick-look" is accomplished through data streaming, which is capable of automatically updating the viewer's web page within seconds.
- The other is a series of increasingly detailed static web pages which are updated several times an hour. These provide reports on file latency, ground station, and customer distribution, and network throughput.

Through use of these reports, ground station personnel, customers, project managers, and system administrators can remotely verify file availability and system performance.

4.0 Conclusion

The benefits derived by NASA's use of this system to date include the following:

- Increased data distribution in a more flexible, reliable, and timely fashion;
- Additional automation of ground stations activities;
- Reduction of interference to automated ground station operations by non-scheduled customer data requests;
- No new firewall rules needed when new projects are added;
- Reduction of interference to customer systems from increasingly automated ground stations;
- Reduction of costs through unattended operation of the SAFS system;
- Ability to do remote diagnosis of problems;
- Standardized user interface for mission data acquisition;
- Reduction of development and maintenance costs due to use of COTS products;
- Central point for customer interaction;
- No new development needed when new standard projects are added.

As we learn lessons from this successful first application of the SAFS system, its value to the science data community should increase as more opportunities present themselves for its use.